

What is claimed is:

1. A method for monitoring a query during runtime, said query involving a plurality of join operations; the method comprising the steps of:

5 running the query according to a first join order; and
 concurrent with running the query, collecting performance statistics about each of the join operations.

2. The method according to claim 1, further comprising the step of:

10 changing the first join order, during running of the query, to a second join order based on the statistics.

3. The method according to claim 2, further comprising the step of:

15 collecting additional statistics about each of the join operations after the first join order is changed to the second join order.

4. The method according to claim 3, further comprising the step of:

20 changing the second join order to either the first join order or a third join order based on the additional statistics.

5. The method according to claim 1, wherein the plurality of join operations include:

a first join that includes a join-from file and a first join-to file;

a second join that includes the join-from file and a second join-to file.

5 6. The method according to claim 5, further comprising the steps of:

determining respective fan-in statistics for the first join and second join; and

changing the first join order to a second join order if the respective fan-in statistics
indicate that the second join is more likely to cause fan-in than the first join.

10 7. The method according to claim 5, further comprising the step of:

determining respective fan-out statistics for the first join and the second join; and

changing the first join order to a second join order if the respective fan-out
statistics indicate that the second join is less likely to cause fan-out than the first join.

15 8. The method according to claim 5, further comprising the steps of:

determining respective fan-in statistics for the first join and second join;

determining respective fan-out statistics for the first join and the second join; and

changing the first join order to a second join order based on a combination of the
respective fan-in and fan-out statistics.

20

9. The method according to claim 1, comprising the steps of:

identifying a predetermined sample size;

performing the step of collecting statistics for the predetermined sample size;

evaluating the collected statistics; and

5 changing the first join order to a second join order based on the collected statistics.

10. The method according to claim 9, comprising the steps of:

collecting additional statistics for substantially all of the query;

10 comparing the additional statistics with the collected statistics; and

adjusting the predetermined sample size, for use by a subsequent query, according to results of the comparing step.

11. The method according to claim 1, further comprising the steps of:

15 running an other query after the query; and

selecting an initial join order for the other query based on the collected performance statistics.

12. A method for optimizing a query join order during runtime, said query involving a plurality of join operations, the method comprising the steps of:

running the query according to a first join order; and

concurrent with running the query, collecting statistics about each of the join

5 operations; and

based on the collected statistics, selecting a preferred join order, while running the query, such that the query continues to run according to the preferred join order.

13. The method according to claim 12, further comprising the steps of:

10 determining respective fan-in statistics for each of the join operations; and

selecting the preferred join order based on the fan-in statistics.

14. The method according to claim 12, further comprising the steps of:

determining respective fan-out statistics for each of the join operations; and

15 selecting the preferred join order based on the fan-out statistics.

15. The method according to claim 12, further comprising the step of:

performing the step of selecting a preferred join order after collecting statistics for a predetermined number of records from a table involved in the query.

20

16. An apparatus for executing a query comprising:

at least one processor;

a memory coupled with the at least one processor; and

a database engine residing in the memory and executed by the at least one

5 processor, the database engine configured to run a query involving a plurality of join operations according to a first join order; and, concurrent with running the query, collect statistics about each of the join operations.

17. The apparatus according to claim 16, wherein the database engine is further

10 configured to select a preferred join order based on the statistics.

18. The apparatus according to claim 17, wherein the database engine further comprises:

a manager configured to determine the select the preferred join order;

an execution engine, coupled with the manager, and configured to execute the

15 query according to the preferred join order; and

a statistics collector, coupled with the manager and the execution engine,

configured to monitor execution of each of the plurality of join operations, capture respective performance data; and communicate the respective performance data to the manager.

20

19. The apparatus according to claim 16, wherein the statistics include respective fan-in statistics and respective fan-out statistics for each of the join operations.

20. The apparatus according to claim 16, wherein the statistics are collected for a predetermined number of records from a table involved in the query.

21. A program product, comprising:

program code configured upon execution to perform the steps of:

running the query according to a first join order;

concurrent with running the query, collecting statistics about each of the join operations, and

based on the collected statistics, selecting a preferred join order, while running the query, such that the query continues to run according to the preferred join order; and

a signal bearing medium bearing the program code.

22. The program product according to claim 21, wherein the program code is further configured to:

determining respective fan-in statistics for each of the join operations; and

selecting the preferred join order based on the fan-in statistics.

23. The program product according to claim 21, wherein the program code is further configured to:

determining respective fan-out statistics for each of the join operations; and

selecting the preferred join order based on the fan-out statistics.